

FAST FAILOVER WITH MULTIPLE SECONDARY NODES

Angshuman Bezbaruah
Anand A. Kekre
Niranjan S. Pendharkar

ABSTRACT OF THE DISCLOSURE

A method, system, computer system, and computer program product to quickly synchronize replicas of data stored at multiple secondary nodes upon failure of a primary node. Secondary nodes to which an update to data is sent are identified. The secondary nodes insert the update into a respective log of updates to a respective copy of the data. When all of the secondary nodes have acknowledged the update, a notification is sent to each of the secondary nodes. In response to receiving the notification, each of the secondary nodes clears the update from the respective log. Data at one secondary node can be synchronized with data at another secondary node when one of the nodes has received updates from the primary node that the other secondary node has not yet received.